

Market Infrastructure, Volatility Dynamics, and Investor Protection in India's Evolving Capital Markets: An Integrated Empirical Assessment

Nirav Nishit, Poojari Pritish Satish, Chinmaya Bhat

Independent Research

DOI: <https://doi.org/10.51583/IJLTEMAS.2026.150900123>

Received: 25 September 2026; Accepted: 30 September 2026; Published: 19 October 2026

ABSTRACT

India's stock exchanges have grown from regional trading floors into some of the busiest markets in the world. Earlier work by the authors examined retail losses and regulatory reform; this paper turns to the infrastructure underneath. It uses the initial public offering of the National Stock Exchange (NSE), completed on 24 September 2026 after a ten-year delay linked to the co-location case, to examine a tension in exchange governance: a for-profit, listed exchange that also supervises the brokers who trade on it. Because the listing is days old, the paper treats it as a pre-listing baseline and sets out a post-listing research agenda. It then reviews broking, margining and corporate-action processing, and two fast-growing segments, Real Estate Investment Trusts (REITs) and mandatory sustainability reporting. Its quantitative core re-estimates asymmetric GARCH models on 4,553 daily Nifty returns from 2007 to 2026. Persistence is 0.984, a half-life of about 42 trading days, against 0.991 and 79 days in a published benchmark, and falls to 0.930 and about 10 days once nine variance breaks are allowed for. A fat-tailed Monte Carlo simulation puts the 99 percent one-day value-at-risk 7 percent, and expected shortfall 16 percent, above the normal approximation, and in a walk-forward backtest over 3,053 days only the fat-tailed GARCH model passes the coverage and independence tests at both the 95 and 99 percent levels. The NSE's investor compensation ceiling of ₹35 lakh is lower in dollars than the limits of the United States' SIPC and the United Kingdom's FSCS but higher as a multiple of GDP per head. Infrastructure has advanced quickly, while compensation ceilings have lagged participation and exchange governance remains less settled. Every figure comes from a named source or is calculated from one, and the qualitative ratings are marked as the authors' judgement.

Keywords: exchange governance, market microstructure, GARCH volatility, ESG disclosure, investor protection

INTRODUCTION

India's equity and derivatives markets have grown from a modest base in the early 1990s into one of the world's largest trading venues by contract volume. Earlier work by the authors looked at that growth from the side of the retail traders who now dominate it, documenting persistent losses in futures and options and assessing the regulatory response. This paper asks a different question. How well is the machinery beneath all that trading governed, measured and protected against failure?

The question has a concrete starting point. On 24 September 2026 the National Stock Exchange of India (NSE) completed its own initial public offering, listing on the Bombay Stock Exchange (BSE) rather than on itself, at an implied valuation of about ₹4.42 lakh crore (CAclubindia, 2026). The NSE had sought a listing since 2016. A governance dispute, the co-location and dark-fibre case, held it up for nearly a decade; the case turned on whether every trading member had fair and equal access to the exchange's own infrastructure (Multibagg AI, 2026). The episode shows a tension that the exchange-governance literature has long described. A demutualised, for-profit exchange answers to shareholders, yet it also supervises brokers who are often its shareholders and clients (Ferrell, 2007).

Three other parts of market infrastructure get less attention in the retail-loss literature, and the paper covers each. The first is market microstructure, meaning the broking, margining and corporate-action systems that decide how a trade is cleared, financed and settled. The second is risk measurement, in particular how persistent volatility is in India's benchmark index; the paper re-estimates asymmetric GARCH models on raw Nifty data, tests how far their persistence depends on regime shifts, and carries the preferred model into a Monte Carlo value-at-risk exercise and a walk-forward backtest. The third is investor compensation, which the paper compares with the schemes in the United States and United Kingdom.

Three questions organise the analysis. How do India's ownership and supervisory safeguards for exchanges compare with those of other countries that let exchanges become for-profit or listed companies? How persistent is volatility in India's benchmark index, and what follows for margining and portfolio risk? And how does India's investor compensation compare, in size and design, with the American and British schemes? Two further segments, Real Estate Investment Trusts (REITs) and mandatory environmental, social and governance (ESG) disclosure, serve as cases where Indian institutions grew quickly under newly written frameworks. The contribution is the integration: strands usually studied separately are examined together, using regulatory filings, exchange disclosures, offer documents and peer-reviewed volatility research, together with the authors' own estimates from raw index data. The literature review comes first, followed by the methodology. The remaining sections cover exchange governance and the NSE listing, market microstructure, REITs and ESG disclosure, the GARCH analysis and its robustness checks, portfolio risk by Monte Carlo simulation, a walk-forward backtest, and the compensation comparison, before the discussion, limitations, conclusion and data availability statement.

LITERATURE REVIEW

Exchange Demutualisation and Self-Listing

Demutualisation is the shift of stock exchanges from member-owned, not-for-profit associations to shareholder-owned corporations. It has been studied widely since the 1990s, when the Stockholm (1993), Helsinki (1995) and Australian (1998) exchanges converted (Ferrell, 2007). Steil (2002) states the central problem. A demutualised exchange wants to maximise trading and listing revenue, but it must also police the members who generate that revenue. The conflict sharpens once the exchange lists its own shares, because its shareholders then have a financial stake in the exchange's decisions as a regulator.

Ferrell (2007) traces the pace of this global wave. He notes that the NSE was created as a demutualised exchange in 1995 instead of converting from a mutual structure, while the BSE demutualised in 2005. He also lists the safeguards that other jurisdictions attach to for-profit and self-listed exchanges. These include ownership caps for any single holder, often 5 percent, as in Singapore and the Philippines; independent-director requirements; and, in Australia, Hong Kong, Singapore and Sweden, supervision of a self-listed exchange by a government regulator and not by the exchange. The Australian Securities Exchange demutualised on 13 October 1998 and listed its own shares the next day. It is the benchmark for a smooth transition, and it made the change under a statutory 5 percent shareholding cap (Cowan et al., n.d.).

Volatility Clustering and the GARCH Family

Engle (1982) introduced the Autoregressive Conditional Heteroscedasticity (ARCH) model. His insight was that return volatility is not constant; it clusters, so large price moves tend to follow large moves of either sign. Bollerslev (1986) generalised the model into the GARCH form used here, which lets volatility depend on its own past as well as on recent shocks. That captures the persistence of volatility regimes much more economically than ARCH. Glosten, Jagannathan and Runkle (1993) and Nelson (1991) added asymmetric responses to bad news, and both models are used below. For India, Pandey and Kumar (2017) model the persistence of conditional volatility in asset returns.

Other Indian studies find that asymmetric models fit better than a plain GARCH(1,1). Sinha (2007) reports that EGARCH describes Sensex volatility best and GJR-GARCH describes the Nifty best. Suganya,

Padmavathi and Sarantinos (2026) study 2015 to 2025 across the Sensex, the Nifty and twenty-seven sector indices, asking how persistence, asymmetry and lag structure have changed against an earlier benchmark. Most indices show stronger clustering and a sharper reaction to bad news, especially banking, fast-moving consumer goods (FMCG) and energy. One study points the other way. Kumar and Maheswaran (2012) allow

for regime shifts in variance that the data themselves identify. Once they do, the persistence and asymmetry that GARCH-type models estimate for Indian returns fall sharply. Models that ignore regime shifts, they conclude, can make volatility look more persistent than it is. Lamoureux and Lastrapes (1990) had made the same argument for US stocks. This caveat is tested directly below, by estimating the models with and without variance breaks.

Modern Portfolio Theory and Simulation-Based Risk Measurement

Markowitz (1952) set out the mean-variance framework behind modern portfolio theory. A portfolio's risk depends on the variance of each holding and also on how the holdings move together, so spreading money across imperfectly correlated assets can cut risk without cutting expected return in proportion. Monte Carlo simulation takes the idea from a single-period formula to a full simulated distribution of outcomes. It helps most when volatility changes over time, as the GARCH literature says it does, because a simulated GARCH volatility path can feed the simulation directly. The original framework, by contrast, assumes constant variance.

Value-at-Risk (VaR), the measure used in the Monte Carlo discussion below, has become the standard risk figure for regulators and institutions because it turns a loss distribution into one interpretable number. Its known weakness is that it says nothing about how bad losses can be beyond the chosen confidence level. Expected Shortfall was proposed to address this: it averages the losses beyond the VaR cut-off, so it reflects how severe the tail is and not only where it starts. Acerbi and Tasche (2002) show that it is a coherent risk measure, and the Basel Committee's market-risk standard uses it at the 97.5 percent level in place of 99 percent VaR (Basel Committee on Banking Supervision, 2019). The paper reports both measures.

Backtesting and Overfitting

Out-of-sample evaluation matters as much as model choice. Bailey, Borwein, López de Prado and Zhu (2014) show that backtest overfitting is widespread, and that the number of strategy variants tried largely determines how misleading a backtest is. Machine-learning forecasters such as Long Short-Term Memory (LSTM) networks (Hochreiter & Schmidhuber, 1997; Fischer & Krauss, 2018) raise the same issue, because they offer many tunable variants. That is why this paper evaluates its risk model in a walk-forward design, in which every forecast uses only data available at the time, and leaves a machine-learning extension to future work.

Comparative Investor Protection

Investor compensation schemes that protect clients when a broker defaults are standard in developed markets, but their design, funding and limits vary. This paper compares India's Investor Protection Fund with the United States' Securities Investor Protection Corporation and the United Kingdom's Financial Services Compensation Scheme. It uses each scheme's own published limits and places India's August 2024 increase beside both. Because the limits are set in different currencies and legal units, it also scales each ceiling by income per head.

METHODOLOGY

The paper combines document analysis with an econometric and simulation exercise. Evidence on governance, listing events and ownership rules comes from regulatory summaries, the NSE's offer documents, exchange disclosures and financial press, each cited where used. Comparative evidence on demutualisation and safeguards abroad comes from Ferrell (2007) and Cowan et al. (n.d.). For volatility, the paper estimates GARCH-family models by maximum likelihood on 4,553 daily log returns of the Nifty 50 from 18 September

2007 to 13 April 2026, computed from closing levels in a public dataset (kalilurrahman, 2026). The data were checked for duplicate dates, missing values and non-positive levels, none of which occur, and the three largest absolute daily moves, 16.3, 13.9 and 13.0 percent, fall on 18 May 2009, 23 March 2020 and 24 October 2008. Four specifications are compared: GARCH(1,1) with normal shocks, GARCH(1,1) with Student-t shocks, GJR-GARCH (Glosten et al., 1993) and EGARCH (Nelson, 1991) with t shocks. The asymmetric GARCH

estimates that NYU Stern's V-Lab publishes for the Nifty, fitted to daily data from 2 July 1990 to 17 July 2026, serve as an external benchmark and are not inputs to the simulation. Robustness to structural change is examined with non-overlapping and rolling fixed windows, with a kurtosis-adjusted cumulative-sum-of-squares test for variance breaks (Inclán & Tiao, 1994; Sansó et al., 2004) followed by estimation with break dummies (Lamoureux & Lastrapes, 1990), and with a two-state Markov-switching variance model (Hamilton, 1989). Portfolio risk is estimated by Monte Carlo simulation of the preferred model and evaluated in a walk-forward backtest with the Kupiec (1995) and Christoffersen (1998) tests. Compensation evidence uses each scheme's published limit, converted at exchange rates on stated dates and scaled by GDP per head.

Three rules govern the use of evidence. Every quantitative figure is traced to a named source or to the authors' own estimation, and where sources differ in emphasis both are reported. Derived figures, such as half-lives, decay shares, VaR values and currency ratios, are recomputed from the reported inputs and shown with them. Judgements that go beyond the sources, such as the summary ratings in the discussion, are labelled as the authors' own. The design is descriptive and integrative, with an econometric component. It does not test causal hypotheses, and its findings assess the current evidence and are not estimates of causal effects.

Several sources are live pages or fast-moving press reports, so each figure is dated to the moment it describes: the V-Lab benchmark of 20 July 2026, the Nifty sample ending 13 April 2026, the NSE Investor Protection Fund corpus as on 31 May 2026, the exchange rates on the dates stated in the investor-protection section, and the listing on 24 September 2026. Later revisions to those pages fall outside the scope of the paper.

Exchange Governance and the Economics of Self-Listing

The NSE's Decade-Long Path to Listing

The NSE first filed draft offer documents with the Securities and Exchange Board of India (SEBI) in December 2016, aiming to raise about ₹10,000 crore (Multibagg AI, 2026). SEBI returned the filing in 2019 because two proceedings remained open: the co-location case, in which some members were alleged to have gained preferential access to the exchange's trading servers, and a related dark-fibre matter (CAclubindia, 2026). The NSE later agreed a revised settlement of about ₹1,491.21 crore, roughly ₹1,224 crore for co-location and ₹268 crore for dark fibre. SEBI gave in-principle approval in 2026, and the Supreme Court dismissed the regulator's remaining case in early September 2026 (Analytics Insight, 2026; Multibagg AI, 2026). The NSE filed its Draft Red Herring Prospectus on 17 June 2026 and opened its offer on 17 September. The offer, priced in a band of ₹1,700 to ₹1,785 a share, was worth about ₹22,569 crore, and the shares listed on the BSE on 24 September 2026. It was the second-largest offer in Indian history, behind Hyundai Motor India's ₹27,870 crore offer in 2024 (CAclubindia, 2026; India TV News, 2026). The Red Herring Prospectus dated 10 September 2026 states that the NSE has no identifiable promoter, and describes the offer as an offer for sale of up to 126,436,650 equity shares of ₹1 face value by existing shareholders, with the BSE as the designated stock exchange (National Stock Exchange of India, 2026b). At the top of the price band that is $126,436,650 \times ₹1,785$, or about ₹22,569 crore, which matches the reported offer size (authors' calculation).

The Regulatory-Conflict Debate

The listing puts the demutualisation literature's central worry in sharp focus. As a listed company, the NSE answers to public shareholders who want more trading volume, listing revenue and data income. Those are the lines in which it already dominates. In FY26 its market shares were about 93 percent in cash equities, 99.79 percent in equity futures and 74.71 percent in equity options (Groww, 2026, citing a Redseer report that, according to the offer document, the NSE commissioned and paid for; National Stock Exchange of India,

2026b). It also remains a first-level supervisor of the members who generate that revenue, responsible for surveillance, margin monitoring and discipline. SEBI's co-location proceedings were about exactly this dual role. Would a listed NSE, facing quarterly earnings scrutiny, invest more or less in costly, revenue-neutral surveillance than an unlisted NSE did? That is an empirical question, and the listing took place only days before this paper was finalised, so there is no post-listing behaviour to observe. The paper therefore treats the listing as a pre-listing baseline. The final subsection of this section records what can be measured on the eve of listing and sets out how a post-listing study could test the question.

The financial backdrop is worth recording. The NSE reported a profit of about ₹10,300 crore for FY26, down 15 percent on the year before (Analytics Insight, 2026), and it paid a ₹6.04 crore penalty in May 2026 over a technical glitch in 2024 (CAclubindia, 2026). Part of the profit decline may reflect lower derivatives activity after the 2024-25 reforms the authors examined earlier, though the sources consulted do not break the decline down. If that is right, any future clash between shareholder returns and further tightening of the derivatives segment now plays out inside a listed company's boardroom, not only between an unlisted exchange and its regulator.

NSE's listing in Comparative Context

Set against the global demutualisation wave, the NSE is unusual. Most exchanges began as member-owned associations and later converted to shareholder-owned corporations. The NSE was founded already demutualised in 1995, so trading rights were separate from ownership from its first day (Ferrell, 2007). The BSE took the conventional route and demutualised in 2005 (Ferrell, 2007). Table 1 sets the NSE's 2026 listing beside the gap between demutualisation and listing at several other major exchanges.

Table 1: Demutualisation-to-Listing Gap at Major Global Exchanges

Exchange	Demutualised	Listed	Gap
Australian Securities Exchange	13-Oct-98	14-Oct-98	1 day
New York Stock Exchange	2006	2006	Same year
London Stock Exchange	2000	2001	~1 year
Nasdaq	2000	2002 (OTC); 2005 (own market)	2 to 5 years
National Stock Exchange of India	1995 (founded demutualised)	2026 (on the BSE)	~31 years

The NSE's gap of about 31 years between demutualised status and public listing is far longer than at any comparator in Table 1, most of which listed within a few years (Ferrell, 2007; Cowan et al., n.d.). It is also the only case in which the shares listed on a different exchange, the BSE, rather than on the issuer itself, which avoids the self-listing conflicts Ferrell (2007) describes. Age and market position do not explain the gap. The NSE was reported to be the world's largest derivatives exchange by trading volume while it remained unlisted (Multibagg AI, 2026). The governance dispute behind the co-location case explains it better, and it illustrates the literature's central worry: a demutualised exchange still answers, as a supervisor, to standards that its own commercial incentives might wear down.

Ownership Limits and Supervisory Safeguards

India handles these conflicts mainly through ownership and board rules in SEBI's Securities Contracts (Regulation) (Stock Exchanges and Clearing Corporations) Regulations, 2018, the SECC Regulations. An individual investor is limited to 5 percent of a stock exchange, and specified institutions to 15 percent. Trading members together with their associates may hold no more than 49 percent, and holdings above 45 percent

need prior approval for further purchases. Shareholders with 2 percent or more must pass fit-and-proper tests (Business Standard, 2024). A board meeting is quorate only if public interest directors at least match shareholder directors (AZB & Partners, 2018). A 2024 SEBI framework adds a further layer: each market infrastructure institution must appoint an unaffiliated designated depository to monitor these limits (Business Standard, 2024). The shareholding-disclosure duties in the same framework apply to these institutions

whether or not they are listed (Business Standard, 2024). Because holders of 2 percent or more must already pass fit-and-proper tests, the framework gives SEBI a way to screen owners well before they near the 5 percent cap. Each rule targets a specific conflict. The 5 percent cap limits any single owner's power to steer the exchange, and the 49 percent ceiling on trading members stops the exchange's main users from controlling it as a group. The quorum rule works on the boardroom itself, since shareholder directors cannot form a quorum without independent voices. The NSE's statement that it has no identifiable promoter fits this diversified-ownership design. Table 2 sets these rules beside the safeguards Ferrell (2007) and Cowan et al. (n.d.) document elsewhere; its US row reflects the position and proposals of 2007 and does not cover later rule changes.

Table 2: Ownership and Supervisory Safeguards for Exchanges Across Jurisdictions

Jurisdiction	Ownership limit	Other safeguards	Source
India (SECC Regulations, 2018)	5% per investor; up to 15% for specified institutions; trading members and associates 49% collectively	Quorum needs public interest directors to match or outnumber shareholder directors; a designated depository monitors the limits	AZB & Partners (2018); Business Standard (2024)
Australia (ASX)	5% statutory cap on any single holder	Self-listed exchange supervised by a government regulator	Cowan et al. (n.d.); Ferrell (2007)
Singapore and Philippines	5% single-entity cap, absent a regulatory waiver	Singapore: self-listed exchange supervised by a government regulator	Ferrell (2007)
United States (as of 2007)	SEC proposed a 20% cap on any one-member broker-dealer	Self-regulatory functions in separate subsidiaries; proposed Regulation AL requiring oversight-committee certification and quarterly reports on self-listed securities	Ferrell (2007)

Read against Table 2, India's ownership regime is at least as detailed as the safeguards Ferrell (2007) documents elsewhere in 2007, and its 5 percent general cap matches the statutory cap that accompanied the ASX's own listing. What ownership rules cannot do is guarantee fair access to trading infrastructure. The co-location dispute was about latency and server access, not concentrated shareholding, and it arose at an exchange whose ownership was already spread out. The authors read this as an inference, not a finding: ownership caps deal with one channel of conflict, while operational fairness depends on continuous technical surveillance and audit, where the public record is thinner.

A Pre-Listing Baseline and A Post-Listing Research Agenda

Table 3 sets out five indicators, each with its pre-listing value from the sources used in this section and a design for testing it after listing. The proposed design compares each indicator over the four quarters before and after 24 September 2026, adjusting for market-wide activity; it is the authors' proposal, not a result. Three hypotheses organise it. First, listing raises earnings pressure, so spending on surveillance and enforcement

outcomes should be tracked against profitability. Second, public-shareholder scrutiny should raise disclosure, so the frequency and detail of infrastructure-audit and outage reporting should rise. Third, the NSE's market shares and fee setting should be tracked, because the conflict is most likely to appear where its commercial interest is largest. None can be tested now. Each could be examined once a few quarters of post-listing data exist, using public filings.

Table 3: Pre-Listing Baseline for the NSE and a Proposed Post-Listing Research Agenda (Authors' Proposal)

Indicator	Pre-listing baseline	Proposed post-listing test	Data source
Enforcement and settlements	Settlement of about ₹1,491.21 crore over the co-location and dark-fibre matters (Analytics Insight, 2026; Multibagg AI, 2026)	Number and size of SEBI penalties and settlements per quarter, before and after listing	SEBI orders; NSE disclosures
Profit pressure	FY26 profit of about ₹10,300 crore, down 15 percent (Analytics Insight, 2026)	Profit margin and spending on surveillance and regulatory functions in quarterly results	NSE quarterly and annual results
Market concentration	FY26 shares of 93 percent in cash equities, 99.79 percent in equity futures and 74.71 percent in equity options (Groww, 2026)	Market shares and fee changes in each segment	Exchange monthly statistics; SEBI
Governance structure	No identifiable promoter (National Stock Exchange of India, 2026b); public interest directors at least match shareholder directors (AZB & Partners, 2018)	Board and committee composition, including the public-interest-director share	Stock-exchange filings
Operational resilience	A 2024 technical glitch, penalised at ₹6.04 crore in May 2026 (CAclubindia, 2026)	Frequency and duration of outages and the detail of audit disclosure	Exchange notices; annual reports

Market Microstructure: Broking, Margining, and Corporate Actions

The evolution of Broking

India's broking industry has moved over the past decade from full-service firms charging percentage brokerage to discount brokers charging a flat fee per order. That cut the cost of frequent trading and, the authors have argued earlier, helped drive retail derivatives participation. The NSE reports 132.37 million unique registered investors at 30 June 2026, spread across more than 99 percent of Indian postal codes (Groww, 2026), a reach that a branch-based model would struggle to match. Scale brings an operational-resilience question, and the point here is an inference, not a sourced finding. If order flow is concentrated in a few technology platforms, an outage or risk failure at one could reach a large share of retail investors.

Position Limits and Systemic Risk Monitoring

Position limits are a further safeguard that complements margining. They cap how large one participant's derivatives exposure may become, both in absolute terms and relative to open interest. SEBI's 2024-25 reform package, examined in the authors' earlier work, added intraday position-limit monitoring to close a gap: limits had been checked only at the end of the day, so large positions could be built and unwound inside a session without appearing in any end-of-day snapshot. Position limits work differently from margins. A margin raises the cost of holding a position, while a limit caps its size whatever capital the participant is willing to commit.

In the vocabulary of the earlier capital-constraint versus structural-constraint framework, a limit is the more directly structural intervention.

Margining Frameworks

Margining is the collateral a trader posts to hold a derivatives position, and it has become a significant regulatory lever, as the authors' earlier work on SEBI's 2024-25 reforms showed for options-buyer premium collection. In stylised terms, Indian clearing corporations compute margins with a portfolio framework of the Standard Portfolio Analysis of Risk (SPAN) type. Such a framework estimates the largest one-day loss a portfolio could plausibly suffer by scanning a grid of price and volatility moves, scaled by a Value-at-Risk parameter that the exchange resets periodically against recent volatility. This links margining to the persistence question addressed later. The scanning range is recalibrated on a look-back basis, not continuously, so after a large move margins lag true risk until the next recalibration. That is a general property of look-back margining and not a flaw peculiar to Indian exchanges. But persistence of the kind documented further below, between 0.98 and 0.99 in single-regime estimates, makes the lag matter more in India than in a market where volatility reverts quickly, although the break-adjusted estimate is much lower.

A related mechanism is the peak-margin framework that SEBI introduced for the cash and derivatives segments. Brokers must collect the full applicable margin at any point in the day when a position is held, not only at the close. That closed a gap that had allowed intraday leverage well beyond what end-of-day margining permitted. Together with the expiry-day Extreme Loss Margin surcharge from the earlier work, margining has moved during 2024-25 toward more frequent and more conservative collateral collection, even though the underlying scanning method has not been rebuilt to include a persistence estimate of the kind reported here.

Corporate-Action Processing

Dividends, stock splits, bonus issues and rights issues are processed through India's dematerialised depository system, run by the National Securities Depository Limited and the Central Depository Services Limited. Their combined demat accounts grew from about 39.3 million in March 2019 to 237.7 million by August 2026, a compound annual growth rate above 27 percent; the figures come from depository data compiled in the authors' earlier work. Processing at that scale without many errors is itself a real achievement, yet it gets little attention because it works well enough to stay invisible. The exchange platforms also support capital formation beyond trading. At 30 June 2026 the NSE had 3,005 listed entities, 565 of them small and medium enterprises, and 195 non-profit organisations registered on its Social Stock Exchange; it facilitated ₹20.33 lakh crore of fund mobilisation in FY26 (Groww, 2026). Accurate corporate-action processing matters here because every listed entity's dividends, splits and bonus issues flow through the same depository rails.

Emerging Market Segments: Reits and Esg Disclosure


REITs as an Institutional Real Estate Access Channel

REITs entered Indian markets under SEBI's 2014 REIT Regulations. Embassy Office Parks REIT made the first listing in 2019, followed by Mindspace Business Parks REIT in 2020, Brookfield India Real Estate Trust in 2021, Nexus Select Trust, the first retail-focused REIT, in 2023, and Knowledge Realty Trust on 18 August 2025, now the largest Indian office REIT by gross asset value (CREDAI & ANAROCK, 2025; Business Standard, 2025b). REITs must pay out at least 90 percent of net distributable cash flows, and reported yields range from about 5 to 7.5 percent depending on the source and window (The Statesman, 2025; CREDAI & ANAROCK, 2025). From 1 January 2026 SEBI treats REITs as equity-related instruments (The Statesman, 2025).

Headline growth needs care because a new listing inflates it. Gross assets under management of the four REITs then listed were about ₹1.63 lakh crore at the end of FY 2024-25 (Indian Infrastructure, 2025), and the Indian REITs Association reported about ₹2.35 lakh crore for the five by Q2 FY26 (Business Standard,

2025b), a rise of 44 percent. Knowledge Realty Trust reported a gross asset value of ₹61,999 crore at 31 March 2025 in its offer documents (Yahoo Finance, n.d.). Subtracting it leaves about ₹1.73 lakh crore for the four incumbents, a like-for-like rise of about 6 percent in six months (authors' estimate; the two dates differ, and assets under management and gross asset value may be defined differently, so the estimate is approximate). Distributions give a cleaner like-for-like measure: the four incumbents paid ₹1,641 crore in Q2 FY26, about 19 percent more than a year earlier (Business Standard, 2025b). By Q1 FY27 the Indian REITs Association counted six REITs with gross assets above ₹3.17 lakh crore (Business Today, 2026), but that total again includes new listings. Figure 1 separates incumbent and new-listing assets.

Figure 1: Gross assets under management of listed REITs in India, separating the four incumbent REITs from the 2025 listing of Knowledge Realty Trust. The incumbent figure for Q2 FY26 is the total less Knowledge Realty Trust's gross asset value at 31 March 2025. Source: authors' calculation from Indian Infrastructure (2025), Business Standard (2025b), Yahoo Finance (n.d.) and Business Today (2026).



The segment is growing but remains narrowly held. The five REITs paid ₹2,331 crore to more than 3.3 lakh unitholders in Q2 FY26 (Ghar.tv, 2025), and by Q1 FY27 the six trusts reported 4.85 lakh unitholders and ₹3,136 crore of quarterly distributions (Business Today, 2026). The 4.85 lakh unitholders compare with 132.37 million unique registered investors on the NSE (Groww, 2026), about 0.37 percent (authors' calculation; the two counts are defined differently, so the ratio is only an order-of-magnitude guide).

Mandatory ESG Disclosure under BRSR

SEBI's Business Responsibility and Sustainability Reporting (BRSR) framework came from a May 2021 amendment to Regulation 34(2)(f) of the Listing Obligations and Disclosure Requirements Regulations. It obliges the 1,000 largest listed companies by market capitalisation to file standardised, quantitative ESG disclosures with their annual reports, mandatory from FY 2022-23 (SEBI, 2021). It is organised around the nine principles of the National Guidelines on Responsible Business Conduct, with essential (mandatory) and leadership (voluntary) indicators. A newer layer, BRSR Core, requires independent assurance of a smaller set of high-scrutiny metrics. On 29 March 2023 SEBI's board approved BRSR Core, value-chain disclosure and assurance, to be phased in over time (KPMG, 2023). Assurance reaches all of the top 1,000 in FY 2026-27, the year that began on 1 April 2026, and value-chain disclosure has been voluntary for the top 250 since FY 2025-26 (GreenSutra, 2026; Keslio, 2026). Adoption has outrun the mandate. Over 1,200 companies publish a BRSR against roughly 1,000 required, a figure SEBI leadership has cited (GreenSutra, 2026).

REITs, ESG and the Governance-Protection Tension

The REIT and ESG strands belong in this paper because they show a regulatory design that the exchange-governance and compensation strands lack. Both frameworks fixed their safeguards in advance and in measurable terms: REITs through a mandatory 90 percent payout and equity-style regulation, and BRSR through standardised quantitative indicators with independent assurance for core metrics. Adoption then ran ahead of the rules, with more than 1,200 companies publishing a BRSR against roughly 1,000 required (GreenSutra, 2026), and the number of listed REITs rose from four in early 2025 to six by Q1 FY27 (Business Today, 2026). The governance and compensation strands show the opposite pattern: safeguards that act only after a failure, such as a compensation ceiling, or that depend on continuing discretion, such as audits of exchange infrastructure, have moved more slowly. The REIT segment also shows the sustainability link directly.

All five REITs listed in 2025 hold 5-star GRESB (Global Real Estate Sustainability Benchmark) ratings, draw between 38 and 74 percent of their power from green sources and have set net-zero targets that run from 2030 to the early 2040s (The Statesman, 2025). The pipeline is also growing: Bagmane Prime Office REIT filed draft offer documents with SEBI in late 2025 to raise up to ₹4,000 crore (Business Standard, 2025a) and is among the six REITs reported by Q1 FY27 (Business Today, 2026), and industry analysis sees a ₹67,000 to ₹71,000 crore monetisation opportunity under SEBI's Small and Medium REIT regulations (CREDAI & ANAROCK, 2025).

A Garch Analysis of Indian Equity Index Volatility

Model Specification

Bollerslev's (1986) GARCH(1,1) model makes the conditional variance of returns depend on two things: the previous period's squared shock and the previous period's own conditional variance:

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2$$

In this equation the left side is the conditional variance of returns on day t . The terms ω , α and β are estimated parameters. The term multiplying α is the previous day's squared shock, and the term multiplying β is the previous day's conditional variance. The sum $\alpha + \beta$ measures persistence. A value near one means a volatility spike decays slowly, and a value near zero means volatility returns quickly to its long-run average. The time a shock takes to decay by half follows directly from the persistence:

$$t_{1/2} = \frac{\ln(0.5)}{\ln(\alpha + \beta)}$$

The paper estimates three extensions of the GARCH(1,1) model. Replacing normal shocks with Student- t shocks lets the data choose how heavy the tails are. The GJR-GARCH model of Glosten, Jagannathan and Runkle (1993) adds a response to negative shocks, so that the conditional variance becomes:

$$\sigma_t^2 = \omega + (\alpha + \gamma I_{t-1}) \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2$$

Here I is an indicator equal to one if the previous shock was negative and zero otherwise, and γ is the leverage parameter: a negative shock raises next-day variance by $\alpha + \gamma$ times its square, and a positive shock by α times its square. With a symmetric shock distribution, half of all shocks are negative, so persistence becomes:

$$P = \alpha + \beta + \frac{\gamma}{2}$$

and the half-life of a shock is $\ln(0.5)$ divided by $\ln(P)$, as above. The EGARCH model of Nelson (1991) instead models the logarithm of variance, which keeps variance positive without parameter restrictions; its persistence is its autoregressive coefficient. The V-Lab benchmark below is a related asymmetric specification

whose coefficients cannot be compared directly with those of the GJR model, but whose persistence and half-life can.

Parameter Estimates for the Nifty Index

The paper estimates the four specifications on the daily returns described in the methodology. For comparison, Table 4 reports the benchmark that New York University Stern School of Business's Volatility Laboratory (V-Lab) publishes for the NSE's Nifty index, an asymmetric GARCH fitted to data from 2 July 1990 to 17 July 2026. The benchmark is not re-estimated here and is not an input to the simulation. Table 5 reports the paper's own estimates.

The sample windows differ. The paper's returns cover 18.6 years and 4,553 observations, including the global financial crisis, the 2020 pandemic crash and the period to April 2026; the V-Lab fit covers more than 36 years and well over 8,000 observations. Fitting a single regime to long samples is exactly the situation in which Kumar and Maheswaran (2012) warn that estimated persistence can be overstated, so the paper tests that caveat directly below. V-Lab refits its model as data arrive, so the coefficients in Table 4 are a snapshot dated 17 July 2026. Large t-statistics reflect the very long sample and are not evidence that the specification is correct.

Table 4: V-Lab AGARCH Benchmark Estimates for the Nifty Index (Not Re-estimated; Fitted 2 July 1990 to 17 July 2026)

Parameter	Estimate	t-statistic	Interpretation
ω (constant)	0.0218	15.59	Unconditional variance weight
α (ARCH)	0.129	63.04	Response to squared shocks
β (GARCH)	0.8622	497.8	Volatility persistence
γ (leverage)	0.3522	24.44	Additional response to negative shocks
Persistence ($\alpha + \beta$)	0.991	—	Near-unit-root; shocks decay slowly
Half-life	~79 trading days	—	Time for a shock to decay by half

In the benchmark, all four parameters are highly statistically significant, and V-Lab displays a persistence of 0.991. Adding the unrounded ARCH and GARCH estimates in Table 4 gives $0.1290 + 0.8622 = 0.9912$, which implies a half-life of about 78 trading days, in line with the roughly 79 days V-Lab reports; the rounded 0.991 alone would imply 76.7 days. The positive, significant leverage parameter (0.3522) confirms the asymmetry described in the literature review. That fits the sector-level evidence of growing asymmetry in Suganya et al. (2026) and the asymmetric specifications that Sinha (2007) finds work best for Indian data.

Re-Estimated Models

Table 5 reports the four specifications estimated on the paper's own sample. Three results stand out. First, fat tails matter: Student-t shocks raise the log-likelihood by about 97 points over normal shocks in the same GARCH(1,1) model and lower the half-life from 103 to 65 days, with an estimated 7.2 degrees of freedom. Second, asymmetry matters: the leverage term is significant in both the GJR specification ($\gamma = 0.132$, $t = 6.6$) and the EGARCH specification, and these two models improve on the symmetric GARCH-t by about 101 and 113 points of the Akaike criterion. Third, once both features are included, estimated persistence is lower than in the benchmark: 0.9835 for GJR-t, a half-life of about 42 trading days, and 0.9851 for EGARCH-t, about 46 days. The best information criteria belong to EGARCH-t, but its advantage over GJR-t is about 11 points, and

the paper uses GJR-t for simulation because its persistence and leverage parameters can be read directly alongside the equations above.

Table 5: Re-estimated Volatility Models for the Nifty Index, 18 September 2007 to 13 April 2026 (4,553 Daily Returns)

Model	α	γ	β	Shock df	Persistence	Half-life (days)	AIC
GARCH(1,1), normal	0.0952	—	0.8981	—	0.9933	103	13,216.90
GARCH(1,1), t	0.0794	—	0.91	7.19	0.9894	65	13,024.70
GJR-GARCH, t	0.0124	0.132	0.9052	7.64	0.9835	42	12,923.40
EGARCH, t	0.1406	-0.0922	0.9851	7.61	0.9851	46	12,912.10
V-Lab AGARCH (benchmark)	0.129	0.3522	0.8622	—	0.9912	79	—

In the EGARCH model the coefficients act on the logarithm of variance, so γ is negative when bad news raises volatility, and its persistence is β ; the V-Lab coefficients follow that platform's own parameterisation and are shown for reference only.

Robustness: Fixed Windows and Regime Shifts


A single set of parameters over 18 years assumes that volatility behaves the same way throughout. Three checks relax that assumption; Table 6 and Figure 2 summarise them. First, fixed windows: re-estimating the GJR-t model on four non-overlapping windows of 1,000 days gives persistence between 0.962 and 0.993, with half-lives from 18 to 98 days, and 0.853, a half-life of about 4 days, for the final 553 days; a rolling 1,250-day window gives a median of 0.967 and a range of 0.885 to 0.994 across 27 windows. Second, structural breaks: a kurtosis-adjusted cumulative-sum-of-squares test (Inclán & Tiao, 1994; Sansó et al., 2004) locates nine breaks in unconditional variance, in August 2009, November 2013, December 2014, May 2016, September 2018, February 2020, April 2021, October 2022 and June 2024. Adding a variance dummy for each regime to the GJR-t model, as Lamoureux and Lastrapes (1990) and Kumar and Maheswaran (2012) do, lowers persistence from 0.9835 to 0.9305 and the half-life from about 42 to about 10 days; the GARCH-t estimate falls from 0.9894 to 0.9405, and its half-life from 65 to 11 days. Third, a two-state Markov-switching model (Hamilton, 1989) separates a calm regime with annualised volatility of 12.9 percent and an expected duration of 93 days from a turbulent regime at 37.9 percent that lasts 25 days on average, and it classifies about 20 percent of days as turbulent.

Table 6: Persistence and Half-Life of Nifty Volatility Shocks Under Alternative Treatments of Structural Change

Specification	Sample	Persistence	Half-life (days)
GJR-t, one regime	2007–2026	0.9835	42
GJR-t, nine variance breaks	2007–2026	0.9305	10
GARCH-t, one regime	2007–2026	0.9894	65
GARCH-t, nine variance breaks	2007–2026	0.9405	11
GJR-t, fixed window 1	18 Sep 2007 – 13 Oct 2011	0.993	98
GJR-t, fixed window 2	14 Oct 2011 – 30 Nov 2015	0.9683	21

GJR-t, fixed window 3	1 Dec 2015 – 31 Dec 2019	0.9619	18
GJR-t, fixed window 4	1 Jan 2020 – 11 Jan 2024	0.9652	20
GJR-t, fixed window 5 (553 days)	12 Jan 2024 – 13 Apr 2026	0.8534	4
GJR-t, rolling 1,250-day windows (median of 27)	2012–2026	0.967	21
V-Lab AGARCH benchmark	1990–2026	0.9912	79

Figure 2: Persistence of the GJR-t model in rolling 1,250-day windows, against the full-sample estimate, the break-adjusted estimate and the V-Lab benchmark. Vertical lines mark the nine variance breaks. Source: authors' estimates from Nifty daily closes (kalilurrahman, 2026); benchmark from NYU Stern V-Lab (2026).



The three checks agree on direction. Persistence is highest in the first window, which contains the 2008 crisis and its aftermath, and in the single-regime full-sample fit; it falls once variance is allowed to shift between regimes, as Kumar and Maheswaran (2012) found for earlier Indian data, and it is lowest in the most recent window. The V-Lab half-life of 79 days is therefore best read as the upper end of a range running from about 10 days, with variance breaks, to about 42 days under a single fat-tailed regime. Even the shortest estimate means that shocks decay over weeks, not days, so the margining implications below survive; what changes is the recalibration interval they imply.


Implications of Persistence for Margin Recalibration

Table 7 turns the decay paths into practical horizons for the three persistence values. One month after a volatility shock, about 83 percent of its initial effect on conditional variance remains under the benchmark, 71 percent under the re-estimated GJR-t model, and 22 percent when variance breaks are allowed for. After a quarter the shares are 57, 35 and 1 percent. Suppose, purely for illustration, that a clearing corporation reset its margin parameters once a month. On the benchmark decay path most of a shock's effect on variance would still be present at the reset, so the reset would need to keep, not discard, most of the elevated risk estimate; on the break-adjusted path less than a quarter would remain, and a monthly reset would be enough. Figure 3 plots the three decay paths.

Table 7: Share of an Initial Volatility Shock Remaining After Selected Horizons Under Three Persistence Estimates (Percent)

Horizon	Trading days	Benchmark (0.9912)	GJR-t (0.9835)	With breaks (0.9305)
1 week	5	95.7	92	69.7
1 month	21	83.1	70.6	22
1 quarter	63	57.3	35.1	1.1
6 months	126	32.8	12.4	<0.1
1 year	252	10.8	1.5	<0.1

Figure 3: Volatility shock decay in the Nifty index under the V-Lab benchmark persistence (0.9912), the re-estimated GJR-t persistence (0.9835) and the break-adjusted GJR-t persistence (0.9305). Source: authors' calculation from NYU Stern V-Lab (2026) and the authors' estimates.



Two cautions apply to reading these figures as forecasts. The estimates rest on parametric specifications fitted to closing-level data, and the break dates are located by a test and then treated as known, which flatters the break-adjusted model; a full Markov-switching GARCH would estimate regimes and persistence together. And the right persistence for margining is not obvious: a margin system that reassesses collateral infrequently, or that relies on a short look-back window, lags a risk environment that normalises slowly after a large move, and the lag is longer the nearer persistence is to the single-regime estimates. That is a design issue separate from, but complementary to, the specific SEBI margin measures examined in the authors' earlier work.

Portfolio Risk and Monte Carlo Simulation

Mean-Variance Foundations

Markowitz's (1952) mean-variance framework expresses portfolio variance as a function of individual asset variances and their pairwise covariances:

$$\sigma_p^2 = \sum_{i,j=1}^N w_i w_j \sigma_i \sigma_j \rho_{ij}$$

Here N is the number of assets, the w terms are portfolio weights on assets i and j , the σ terms are their individual volatilities, and ρ is the correlation between them, equal to one for an asset with itself. For a GARCH-aware framework the key point is that the volatilities, and possibly the correlations, vary over time under the volatility model of the previous section. The original closed-form framework assumes they are constant. The formula also shows why adding an asset that moves against the rest of the portfolio can lower total risk even when that asset is itself volatile.

A two-asset example makes the equation concrete. Take equal weights of 0.5 in two assets that each have 20 percent annual volatility. If the assets are perfectly correlated, portfolio volatility is 20 percent and diversification does nothing. At a correlation of 0.5 it falls to about 17.3 percent, at zero correlation to about 14.1 percent, and at a correlation of minus 0.5 to 10 percent. The numbers are hypothetical, but they show why the correlation term matters as much as the individual volatilities. They also show why any tendency for correlations to rise in stress, which a constant-parameter formula cannot capture, would erode the diversification benefit when it is needed most.

Monte Carlo Value-At-Risk

This paper implements the Monte Carlo simulation instead of only describing it. The engine is the GJR-GARCH model with Student- t shocks estimated above ($\omega = 0.0197$, $\alpha = 0.0124$, $\gamma = 0.1320$, $\beta = 0.9052$, 7.64 degrees of freedom). On each of 100,000 paths the shock on each day is a standardised Student- t draw scaled by the day's conditional volatility, and the next day's variance follows the equation above, so volatility clusters and rises after falls.

Paths run for one or ten trading days, mean returns are set to zero as in the parametric calculation, and the loss on a hypothetical ₹10,00,000 portfolio tracking the Nifty is $V(1 - e^R)$, where R is the cumulative log return. VaR at level α is the α quantile of simulated losses. Expected Shortfall, the average loss at or beyond that quantile, is:

$$ES_{\alpha} = \frac{1}{1 - \alpha} \int_{\alpha}^1 VaR_u du$$

Results are averaged over five random seeds. Across seeds the standard deviation of the 99 percent one-day VaR is about ₹124, or 0.6 percent of its value, and that of the 99 percent one-day Expected Shortfall about ₹221, so the tail estimates are stable.

Results: Normal Approximation against Monte Carlo

The benchmark is a simplified parametric calculation. NYU Stern V-Lab's 1-day-ahead AGARCH forecast for the Nifty, published for 21 July 2026, was 13.76 percent annualised (NYU Stern V-Lab, 2026). Dividing the annualised figure by the square root of the 252 trading days in a year gives a daily standard deviation.

One-day Value-at-Risk at confidence level α is the matching normal quantile multiplied by that daily volatility and the portfolio value V . The simulation starts from the same 13.76 percent forecast, so that the comparison isolates the effect of the return distribution and the volatility dynamics. The formulas are:

$$\sigma_d = \frac{\sigma_A}{\sqrt{252}}$$

$$VaR_{\alpha} = z_{\alpha} \times \sigma_d \times V$$

With an annualised volatility of 13.76 percent, the daily standard deviation is about 0.867 percent. For a hypothetical ₹10,00,000 portfolio tracking the Nifty, the normal-approximation one-day VaR is about ₹14,258 at 95 percent and ₹20,165 at 99 percent. Table 8 compares these figures with the Monte Carlo results, using square-root-of-time scaling for the ten-day normal figures.

Table 8: One-Day and Ten-Day Value-at-Risk and Expected Shortfall for a ₹10,00,000 Nifty Portfolio, Normal Approximation Against GJR-GARCH-t Monte Carlo (Starting Volatility 13.76 Percent a Year)

Horizon	Confidence	Normal VaR (₹)	Monte Carlo VaR (₹)	Monte Carlo ES (₹)	VaR difference
1 day	95%	14,258	13,874	18,747	-2.7%
1 day	97.50%	16,989	17,155	22,145	1.00%
1 day	99%	20,165	21,487	26,859	6.60%
10 days	95%	45,086	46,156	63,674	2.40%
10 days	97.50%	53,724	58,007	75,902	8.00%
10 days	99%	63,767	73,500	92,755	15.30%

The fat-tailed simulation does not lift every figure. At 95 percent the one-day VaR is ₹13,874, about 3 percent below the normal figure, because a Student-t distribution scaled to the same variance has less mass in its shoulders: its 95 percent quantile is 1.61 standard deviations, against 1.645 for the normal. The ordering reverses further out. The 99 percent VaR is ₹21,487, 7 percent above the normal figure, and the 99 percent Expected Shortfall is ₹26,859, 16 percent above the normal figure of ₹23,102. Thickening the tails at equal variance therefore raises the tail measures but can slightly lower a mid-tail quantile, so a normal approximation is not simply too low at every level.

Over ten days the gap is wider and has a different source. The 99 percent ten-day VaR is ₹73,500 against ₹63,767 from scaling the normal figure, 15 percent more, and Expected Shortfall is ₹92,755 against ₹73,055, 27 percent more. Replacing the t shocks with normal shocks in the same GJR dynamics gives a ten-day 99 percent VaR of ₹73,545, almost identical to the fat-tailed figure, so at this horizon the difference comes from the volatility dynamics and not from the tails: starting at 13.76 percent, below the model's long-run level of about 17.3 percent, variance is expected to rise, and the leverage term adds to losses after falls. Two further points follow. The 97.5 percent Expected Shortfall, which the Basel Committee's market-risk standard uses in place of 99 percent VaR (Basel Committee on Banking Supervision, 2019), is ₹22,145 for one day, 10 percent above the normal 99 percent VaR of ₹20,165; for a normal distribution the two are almost equal (₹20,264 against ₹20,165), so the gap is itself a measure of tail thickness. And the result depends on the starting volatility. Starting instead from the model's own forecast at the end of the sample, 24.18 percent annualised after a turbulent March and April 2026, when the standard deviation of the last 20 daily returns was 1.76 percent, gives a one-day 99 percent VaR of ₹37,449 against ₹35,433 under the normal approximation, and an Expected Shortfall of ₹46,692.

Figure 4: Simulated one-day profit and loss on a ₹10,00,000 Nifty portfolio, GJR-GARCH-t Monte Carlo against a normal distribution with the same volatility (13.76 percent a year). Panel (b) plots the probability that the loss exceeds each threshold. Source: authors' simulation of 200,000 paths.




Figure 4 shows the same result as a distribution. In panel (a) the two curves share a standard deviation, so the difference lies in shape: the simulated distribution is more peaked at the centre and heavier in both tails. Panel (b) plots the probability that the one-day loss exceeds a given amount. The curves cross near the 5 percent line, which is why the 95 percent VaR barely moves, and then separate: at a loss of ₹30,000 the normal curve gives a probability of about 0.03 percent and the simulation about 0.19 percent. A risk manager who relied on the normal curve would therefore be most wrong exactly where the losses are largest.

Walk-Forward Backtest of the Value-At-Risk Models

Design

In walk-forward back-testing a model is re-estimated repeatedly and judged only on the data that immediately follow, which avoids the look-ahead bias that in-sample or shuffled splits can introduce without anyone noticing (Bailey et al., 2014). Here the first 1,500 daily returns form the initial window. Every 21 trading days the GJR-GARCH-t and GARCH-N models are re-estimated on the latest 1,500 returns, and each day's one-day VaR at 95 and 99 percent is computed using only information up to the previous close. That gives 3,053 out-of-sample days, from 6 November 2013 to 13 April 2026. Two simple benchmarks are included: a normal VaR using the standard deviation of the last 250 returns, and historical simulation using the last 500 returns. A VaR model is well calibrated if losses exceed it on a fraction p of days. Kupiec's (1995) test compares the observed number of exceptions x in n days with the expected number, using the statistic:

$$LR = -2 \ln[(1 - p)^{n-x} p^x] + 2 \ln \left[\left(1 - \frac{x}{n}\right)^{n-x} \left(\frac{x}{n}\right)^x \right]$$

which is chi-squared with one degree of freedom if the model is correct. Christoffersen's (1998) independence test asks whether exceptions cluster, and the conditional-coverage statistic adds the two.

RESULTS


Table 9 and Figure 5 report the results. At 95 percent every model has an exception rate within about 0.6 percentage points of the target, and the two models with normal shocks and the GJR-GARCH-t model pass the conditional-coverage test, while the 250-day normal and historical-simulation benchmarks fail it because their exceptions cluster (independence $p = 0.001$ and 0.009 , respectively). The 99 percent level separates the models. The two models with normal shocks, the 250-day normal and GARCH-N, breach the 99 percent VaR on 1.83 and 1.77 percent of days, almost twice the target, and fail Kupiec's test with p -values below 0.001 . Historical simulation has an acceptable exception rate (1.18 percent, $p = 0.33$) but its exceptions cluster ($p < 0.001$), as expected from a method that adapts slowly after volatility rises. Only the GJR-GARCH-t model passes all three tests at both levels, with 164 exceptions at 95 percent (5.37 percent; 153 expected) and 38 at 99 percent (1.24 percent; 31 expected).

Table 9: Walk-Forward VaR Backtest, 3,053 Out-of-Sample Days (6 November 2013 to 13 April 2026); Expected Exceptions Are 153 at the 95 Percent Level and 31 at the 99 Percent Level

Model	VaR level	Exceptions (rate)	Kupiec p	Independence p	Conditional coverage p
Normal, 250-day σ	95%	144 (4.72%)	0.47	0.001	0.004
Normal, 250-day σ	99%	56 (1.83%)	<0.001	0.02	<0.001
GARCH-N	95%	172 (5.63%)	0.11	0.17	0.11
GARCH-N	99%	54 (1.77%)	<0.001	0.96	<0.001
GJR-GARCH-t	95%	164 (5.37%)	0.35	0.68	0.59
GJR-GARCH-t	99%	38 (1.24%)	0.19	0.33	0.26

Historical (500 days)	95%	157 (5.14%)	0.72	0.009	0.03
Historical (500 days)	99%	36 (1.18%)	0.33	<0.001	0.002

Figure 5: Walk-forward one-day 99 percent VaR. Panel (a) shows daily returns, the GJR-GARCH-t and 250-day normal VaR lines, and the GJR-GARCH-t exceptions; panel (b) shows each model's exception rate against the 1 percent target, with the Kupiec p-value. Source: authors' calculation from Nifty daily closes (kalilurrahman, 2026).



Structural breaks and what the Backtest Does Not Show

The out-of-sample window spans the October 2024 to April 2025 derivatives reforms, which changed contract sizes, expiry frequency and margin requirements, so a model trained across that break could mistake a regulatory change for predictive skill. Re-estimating on a rolling window limits that risk, because each forecast uses only the latest 1,500 days, and the model keeps its coverage in recent years: from January 2021 the GJR-GARCH-t model records 76 exceptions in 1,303 days at 95 percent (5.83 percent) and 15 at 99 percent (1.15 percent). Three limits apply. The backtest evaluates VaR and not Expected Shortfall, for which backtest need more care. It covers one index and a unit-beta portfolio. And a model can seem to forecast volatility well simply by learning that today looks like yesterday, a signal anyone can get for free, which is why the comparison includes the simple benchmarks. A hybrid in which a machine-learning volatility forecast, such as an LSTM network, feeds the same VaR engine is a natural extension and would need the same walk-forward protocol.

International Comparison of Investor Protection Frameworks

India's Investor Protection Fund (IPF) is run separately by the NSE and the BSE. It compensates investors when a trading member is declared a defaulter or expelled, up to a fixed ceiling per investor per defaulting member. The NSE raised its ceiling from ₹25 lakh to ₹35 lakh for claims arising from members declared defaulters or expelled after 13 August 2024. Its IPF Trust corpus was ₹2,912.73 crore on 31 May 2026, of which ₹1,779.06 crore related to the cash market and ₹1,026.22 crore to futures and options (National Stock Exchange of India, 2026a). The BSE's equity-segment ceiling is ₹16 lakh for members declared defaulters after 29 May 2024 (Share.market, 2026). The comparison below sets these figures beside the United States' Securities Investor Protection Corporation (SIPC) and the United Kingdom's Financial Services Compensation Scheme (FSCS); Tables 10 and 11 give the detail.

The comparison needs care, so the paper makes it in three ways; Table 10 reports the first two. First, in dollars, at exchange rates on the date of the NSE increase. On 14 August 2024, the first business day after the new ceiling took effect, the rupee closed at ₹83.94 per US dollar (News On AIR, 2024) and sterling stood at

about ₹107.09 (OFX, 2024), so the NSE's ₹35 lakh ceiling was about US\$41,700, and SIPC's US\$500,000 limit was about 12 times larger (SIPC, 2026). The FSCS limit for investments is £85,000 per person per firm, set in April 2019 (FSCS, 2019), or about US\$108,400 at that date, 2.6 times the NSE ceiling. At late-September 2026 rates (₹95.892 per dollar on 25 September and ₹128.408 per pound on 17 September; CEIC, 2026a, 2026b) the multiples are 13.7 and 3.1. The December 2025 rise in the FSCS limit to £120,000 applied to bank deposits, not investments (FSCS, 2026). Second, relative to income. Dividing each ceiling by GDP per head in 2024 (World Bank, 2025) reverses the ranking: the NSE ceiling is about 15.5 times India's GDP per head of US\$2,695, against 5.9 times for SIPC (US\$84,534) and 2.0 times for the FSCS (US\$53,246). Third, by design: Table 11 sets out what each limit covers, since the schemes differ in trigger, coverage unit and cash treatment (U.S. Securities and Exchange Commission, n.d.), so the ratios in Table 10 are indicative and not exact.

Table 10: Investor Compensation Ceilings in Three Jurisdictions, in Dollars, Relative to the NSE Ceiling and Relative to GDP per Head

Scheme	Ceiling	US\$ at 14 Aug 2024	Multiple of NSE ceiling, 14 Aug 2024	Multiple of NSE ceiling, late Sep 2026	Multiple of GDP per head (2024)
India, NSE IPF	₹35 lakh per investor per defaulting member (from 13 Aug 2024)	41,696	1	1	15.5
India, BSE IPF	₹16 lakh per investor per defaulting member (equity segment)	19,061	0.46	0.46	7.1
United States, SIPC	US\$500,000 per customer, including US\$250,000 for cash	500,000	12	13.7	5.9
United Kingdom, FSCS	£85,000 per eligible person per firm (investments)	108,438	2.6	3.1	2

A second normalisation looks at the fund behind the ceiling. The NSE's IPF Trust corpus of ₹2,912.73 crore could pay about 8,322 claims at the full ₹35 lakh ceiling, or ₹220 for each of the 132.37 million unique registered investors (authors' calculation; the corpus is dated 31 May 2026 and the investor count 30 June 2026). Neither number says how likely a default is, only how much is pooled against it. The direction of recent change is also informative. The NSE raised its ceiling by 40 percent in August 2024, from ₹25 lakh to ₹35 lakh, while registered demat accounts grew from 39.3 million in March 2019 to 237.7 million by August 2026, figures from depository data compiled in the authors' earlier work. Per-investor protection has risen, but by far less than participation. In income-scaled terms the ceiling is not low; what has lagged is the rule for updating it.

Table 11: Conceptual Mapping of Compensation Schemes (Sources: National Stock Exchange of India, 2026a; Share.market, 2026; SIPC, 2026; U.S. Securities and Exchange Commission, n.d.; FSCS, 2019, 2026)

Feature	India (NSE and BSE IPF)	United States (SIPC)	United Kingdom (FSCS)
Coverage unit	Per investor per defaulting member; a separate trust at each exchange	Per customer, per separate legal capacity	Per eligible person per authorised firm
Trigger	Trading member declared a defaulter or expelled	Failure of a member brokerage firm with customer property missing	Failure of an authorised firm
Assets cover	Investor claims on the defaulting member; the BSE figure is for the equity	Securities and cash held for investing, with cash capped at US\$250,000 inside the US\$500,000	Investments; the £120,000 limit raised in December 2025 applies to deposits

ed	segment	limit	
Ceiling	₹35 lakh at NSE (from 13 August 2024); ₹16 lakh at BSE	US\$500,000	£85,000 (set April 2019)
Funding	Exchange-administered trust; NSE corpus ₹2,912.73 crore at 31 May 2026	Assessments on member broker-dealers	Levies on authorised firms

Grievance Redress alongside Compensation

Compensation after a member's default is only one layer of investor protection. The second layer is day-to-day complaint handling. SEBI's SCORES 2.0 platform, live since April 2024, gives regulated entities 21 calendar days to file an action-taken report on a complaint. A dissatisfied investor can ask for a first-level review, usually by the exchange or depository, within 15 days. SEBI itself offers a second review within a further 15 days, after which the case can go to online dispute resolution. Share.market (2026) cites monthly figures of over 5,000 complaints resolved in May and June 2026, with average entity response times of four to five days. For serious misconduct, SEBI can impose fines of up to ₹25 crore, or three times the gains made if that is larger, for insider trading and for fraud or unfair dealing (Sections 15G and 15HA of the SEBI Act; Share.market, 2026). These tools address a different failure from a broker default. The two layers are complements, not substitutes: a compensation ceiling limits the loss from a failed intermediary, while complaint rules shorten the time in which an ordinary dispute can grow into a loss. In the authors' reading, India's protection architecture is thinner on the compensation ceiling than on process, a distinction that the per-claim comparison in Table 10 cannot capture alone.

DISCUSSION

The four strands point to a common theme. India's capital-market infrastructure has advanced fastest where it touches trading volume and product innovation, in index options, REITs and discount broking. It has moved more slowly where it touches loss absorption and operational governance, meaning investor compensation and continuous auditing of exchange infrastructure. The NSE's listing shows the asymmetry clearly. Its commercial success, meaning scale, profit and market dominance, made the listing large and valuable. Its governance problems, surfaced by the co-location case, delayed the same listing for a decade. The persistence estimates sharpen the point. Shocks to Nifty volatility take about 42 trading days to fade by half under a single fat-tailed regime, about 10 days once variance breaks are allowed for, and 79 days in the published benchmark; even at the shortest of these, margining and surveillance systems must respond within weeks. That is exactly the kind of investment that sits in tension with a newly listed exchange's quarterly earnings incentives, a tension that Table 3 turns into a testable agenda because the listing is so recent.

The investor-protection comparison no longer points to a simple gap. In dollars, the NSE's ceiling is the lowest of the three: about US\$41,700 against SIPC's US\$500,000 and the FSCS's US\$108,400. Scaled by income per head, the ranking reverses: the NSE ceiling is 15.5 times India's GDP per head, against 5.9 for SIPC and 2.0 for the FSCS (Table 10). What has lagged is the growth of the ceiling relative to participation, since a 40 percent rise in August 2024 compares with a six-fold rise in demat accounts since 2019. Coverage limits are a policy lever that regulators can move directly, unlike the deeper question of exchange self-listing, and the evidence favours an indexation rule over a one-off rise in the level. Two limits apply: GDP per head is a mean and says little about the balances of investors in listed derivatives, and the schemes differ in design (Table 11), so neither ranking is decisive on its own.

The REIT and ESG findings are the counterpoint to the governance-versus-protection tension, as the section on those segments argues. Both frameworks are young: REITs date from 2019 and mandatory BRSR disclosure from FY 2022-23. Adoption has run ahead of the regulatory floor in both. More than 1,200 companies publish a BRSR against roughly 1,000 required, and once the 2025 listing of Knowledge Realty Trust is removed, REIT assets under management grew about 6 percent between March and September 2025, while distributions by the four incumbents rose about 19 percent year on year. Where India's regulators

specified a measurable framework from the start, participants adopted it quickly. That contrasts with the decade the co-location case took to resolve and the pace at which compensation ceilings are updated.

The distinction between position limits and margining, drawn in the microstructure section, offers a testable lens on SEBI's reform package. A limit caps exposure outright and a margin only raises its cost, so the two should affect participation and risk-taking differently across trader segments. This mirrors the capital-constraint versus structural-constraint distinction that the authors' earlier work develops formally. That work grouped all six 2024-25 measures under the capital-constraint heading. Classifying position limits separately is a natural next step that this paper identifies but does not take.

Summary Assessment

The ownership safeguards in Table 2 add a nuance to the governance discussion. India's mix of ownership caps, trading-member ceilings and public-interest-director quorum rules is detailed by international standards, and the NSE's dispersed ownership follows directly from it. Yet the co-location episode arose under an equally dispersed ownership structure. The remaining governance risk therefore lies less in who owns the exchange than in how continuously and independently the fairness of its infrastructure is audited. Table 12 summarises the assessment across the six dimensions examined. The ratings are the authors' qualitative judgements based on the evidence assembled here, not measured indices. Read across the rows, the two weakest points are also the ones regulators can change most directly: the rule for updating the compensation ceiling, which is a policy parameter, and the disclosure of infrastructure audits, which is a reporting choice. The strongest rows, REITs and ESG disclosure, are the newest frameworks.

Table 12: Summary Assessment of India's Capital-Market Infrastructure (Authors' Qualitative Judgement)

Dimension	Key evidence in this paper	Assessment
Exchange governance	Detailed ownership caps and board-quorum rules; 31-year demutualisation-to-listing gap; co-location dispute	Ownership safeguards substantial; operational-fairness record mixed
Market microstructure	132.37 million registered investors; margins and position limits tightened in 2024-25; look-back margining	Advanced, with recalibration lag exposed by volatility persistence
Volatility and risk modelling	Re-estimated GJR-t persistence 0.984 (half-life about 42 days), 0.930 (about 10 days) with variance breaks; fat-tailed VaR passes the backtest	Strong modelling capacity; persistence estimates sensitive to regime treatment
REITs	Six listed REITs by Q1 FY27; like-for-like asset growth about 6 percent against a 44 percent headline; 90 percent distribution rule	Rapid and well-structured growth
ESG disclosure	More than 1,200 firms report against roughly 1,000 required; BRSR Core assurance phasing in	Adoption ahead of mandate
Investor protection	NSE ₹35 lakh: lowest in US dollars, highest as a multiple of GDP per head (15.5 against 5.9 for SIPC and 2.0 for FSCS)	Ceiling not low relative to income; update rule lags participation

Implications for Regulators and Market Participants

Four implications follow from the evidence, and they are suggestions, not findings. Compensation ceilings could be tied to a stated indexation rule linked to participation or account growth, so that per-investor protection keeps pace with the investor base; Table 10 shows that the case is for a rule, not a one-off rise, because the current ceiling is not low relative to income. Clearing corporations could test persistence-aware margin recalibration, since the recalibration interval it implies runs from weeks, on break-adjusted estimates, to months, on single-regime estimates, and the choice of treatment is itself a policy decision. Exchanges and SEBI could publish independent audits of trading-infrastructure fairness more regularly, because ownership

caps do not cover that channel and the public record on it is thinnest. And the speed with which REITs and BRSR were adopted suggests that standardised, quantitative disclosure templates can be a cheap way to change market behaviour, a lesson that could carry over to infrastructure-resilience reporting.

LIMITATIONS

Several limitations qualify the findings. The GARCH models are estimated by the authors on daily closing levels from a public repository, from September 2007 to April 2026; closing data cannot capture intraday volatility, the sample is shorter than the V-Lab benchmark, and it ends before the benchmark's July 2026 forecast date, so the Monte Carlo is seeded with the published 13.76 percent forecast as a scenario and, as a sensitivity, with the model's own end-of-sample forecast of 24.18 percent. Persistence estimates remain sensitive to regime treatment: variance breaks are located by a test and then treated as known, which is the approach of Kumar and Maheswaran (2012) and Lamoureux and Lastrapes (1990) but not a full Markov-switching GARCH. The Monte Carlo and backtest cover one index and a unit-beta portfolio, assume zero mean returns, and evaluate VaR and not Expected Shortfall. The NSE governance discussion rests on a listing that happened days before the paper was finalised, so the paper offers a baseline and a research agenda, not a test. The demutualisation comparison in Table 1 uses a small, non-random set of well-documented exchanges, not a systematic global sample.

The investor-protection comparison uses limits that are not conceptually equivalent across schemes; Table 11 sets out the differences. Exchange rates are dated as stated, but the late-September 2026 multiples combine a sterling rate of 17 September with a dollar rate of 25 September, and GDP per head is a mean income measure and not a measure of typical investor wealth, so neither ranking is decisive. The descriptions of SPAN-type margining and peak-margin collection in the microstructure section are stylised characterisations of regulatory practice, not readings of specific clearing-corporation circulars. Several REIT, listing and market-share figures still come from financial press, industry bodies or an industry report commissioned by the NSE rather than directly from primary filings, and the like-for-like REIT estimate combines figures with different dates and definitions; Table 13 flags each group. In two places (REIT assets under management and the timing of SEBI's settlement approval) secondary sources differ in emphasis, which is why those figures are reported with their sources. Table 12 is a qualitative synthesis by the authors, and the deliberately broad scope means each strand could be developed into its own study.

CONCLUSION

This paper examined four dimensions of India's capital-market infrastructure beyond the retail-trading focus of the authors' earlier work. On exchange governance, it used the NSE's delayed initial public offering, its 31-year gap between demutualisation and listing, and a comparison of India's ownership safeguards with those of other jurisdictions, and it set out the listing as a pre-listing baseline with a post-listing research agenda. On market microstructure, it covered broking, margining, position limits and corporate-action processing. On risk architecture, it re-estimated asymmetric GARCH models on 4,553 daily Nifty returns, found persistence of 0.984 and a half-life of about 42 trading days under a fat-tailed GJR specification, against 0.991 and 79 days in a published benchmark and 0.930 and about 10 days with variance breaks, and showed through Monte Carlo simulation and a walk-forward backtest that a fat-tailed GARCH model gives better-calibrated value-at-risk than normal-shock alternatives. On investor protection, it compared India with the United States' SIPC and the United Kingdom's FSCS in dollars, at dated exchange rates, and relative to income per head. The REIT and BRSR case studies give a more positive counterpoint, since adoption ran ahead of the mandate in both, and like-for-like REIT growth is far lower than the headline figure.

The recurring finding is an uneven pace of maturity. Trading infrastructure, product range, disclosure adoption and volatility-modelling capacity have advanced rapidly. Investor-compensation ceilings are the lowest of the three schemes in dollar terms but the highest relative to income per head, and their growth has lagged participation. Exchange governance pairs detailed ownership safeguards with an operational-fairness record that the co-location case shows to be less settled. Two qualifications apply: the persistence estimate is sensitive to how regime shifts are treated, and the NSE's post-listing behaviour cannot yet be observed. Future

research could extend the capital-constraint versus structural-constraint classification to position limits, estimate a Markov-switching GARCH that feeds the Monte Carlo engine, test whether a machine-learning volatility forecast improves on GARCH in the same walk-forward backtest, or follow the NSE's post-listing governance using the indicators in Table 3 once a full year of data exists. For policymakers, the most tractable steps are the two levers the paper highlights, an indexation rule for the compensation ceiling and the disclosure of infrastructure audits, since both can be changed directly. For researchers, the most useful next study is the post-listing analysis, because that is the evidence this paper could not yet observe.

Data Availability

The Nifty 50 daily closing levels used for estimation, simulation and backtesting come from a public repository (kalilurrahman, 2026). Every other input is a cited public source, dated as stated in the text. The authors' estimates, simulation outputs and the data behind every table and figure are provided in the accompanying Excel workbook, which recalculates the half-lives, decay shares, parametric VaR, Kupiec statistics, currency conversions, income-scaled multiples and like-for-like REIT estimates from the stated inputs. The estimation code is available from the corresponding author on request. Table 13 classifies each group of figures by source type and by whether the authors calculated it.

Table 13: Data Sources and Calculation Status

Figures	Source type	Status
Nifty 50 daily closes and log returns	Public dataset repository	Authors' calculation (returns)
GARCH, GJR and EGARCH estimates; break tests; Markov-switching; rolling windows	Estimated from the dataset	Authors' estimate
V-Lab AGARCH parameters and 13.76 percent forecast	Third-party analytics platform (NYU Stern)	Sourced; half-lives and decay shares are the authors' calculation
Monte Carlo VaR and ES; walk-forward backtest	Simulated and estimated from the dataset	Authors' calculation
NSE promoter status, offer structure, listing venue	Primary: NSE Red Herring Prospectus	Sourced; offer value cross-check is the authors' calculation
NSE profit, penalties, settlement, market shares	Financial press; an industry report commissioned by the NSE	Sourced (secondary)
Ownership limits and board rules	Law-firm and press summaries of SEBI regulations	Sourced (secondary)
Compensation ceilings, NSE corpus, SIPC and FSCS rules	Primary: NSE, SEC, SIPC and FSCS publications	Sourced; conversions, multiples and corpus ratios are the authors' calculation
Exchange rates	FBIL reference rates via CEIC; All India Radio; OFX	Sourced
GDP per head	World Bank, via FRED	Sourced; multiples are the authors' calculation
REIT and BRSR figures	Industry-body releases, financial press, industry reports	Sourced (secondary); like-for-like REIT estimates are the authors' calculation
Summary ratings in Table 12	Authors' judgement	Not measured

ACKNOWLEDGEMENTS

The authors declare no funding was received for this research and no financial interest or benefit arising from its direct application.

REFERENCES

1. Acerbi, C., & Tasche, D. (2002). On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7), 1487–1503.
2. Analytics Insight. (2026). NSE IPO explained: What delayed India's largest exchange listing for a decade. Analytics Insight. <https://www.analyticsinsight.net/news/nse-ipo-explained-what-delayed-indias-largest-exchange-listing-for-a-decade>
3. AZB & Partners. (2018, December 31). Securities Contracts (Regulation) (Stock Exchanges and Clearing Corporations) Regulations, 2018. AZB & Partners. <https://www.azbpartners.com/bank/securities-contracts-regulation-stock-exchanges-and-clearing-corporations-regulations-2018/>
4. Bailey, D. H., Borwein, J. M., López de Prado, M., & Zhu, Q. J. (2014). Pseudo-mathematics and financial charlatanism: The effects of backtest overfitting on out-of-sample performance. *Notices of the American Mathematical Society*, 61(5), 458–471.
5. Basel Committee on Banking Supervision. (2019). Minimum capital requirements for market risk. Bank for International Settlements. <https://www.bis.org/bcb/publ/d457.htm>
6. Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327. [https://doi.org/10.1016/0304-4076\(86\)90063-1](https://doi.org/10.1016/0304-4076(86)90063-1)
7. Business Standard. (2024, October 14). All market infrastructure institutions, whether listed or unlisted, need to disclose shareholding patterns, says Sebi. Business Standard. https://www.business-standard.com/markets/capital-market-news/all-market-infrastructure-institutions-whether-listed-or-unlisted-need-to-disclose-shareholding-patterns-says-sebi-124101400954_1.html
8. Business Standard. (2025a, December 31). Bagmane Group files draft papers with Sebi for ₹4,000-cr REIT IPO. Business Standard. <https://www.business-standard.com/topic/reit>
9. Business Standard. (2025b, November 20). India's listed REITs distribute ₹2,331 cr in Q2 as AUM rises to ₹2.35 trn. Business Standard. https://www.business-standard.com/markets/news/india-reits-q2-fy26-distributions-aum-market-cap-growth-125112000626_1.html
10. Business Today. (2026, August 18). 6 listed REITs, ₹3.17 lakh crore assets: How big is India's REIT market now? Business Today. <https://www.businesstoday.in/personal-finance/story/6-listed-reits-rs3-17-lakh-crore-assets-how-big-is-indias-reit-market-now-549671-2026-08-18>
11. CAclubindia. (2026). NSE IPO 2026: Everything CA/CS professionals need to know. CAclubindia. <https://www.caclubindia.com/articles/nse-ipo-2026-everything-ca-cs-professionals-need-to-know-56259.asp>
12. CEIC. (2026a). India foreign exchange rate: RBI reference rate: US dollars [FBIL reference rate, 25 September 2026]. CEIC Data. <https://www.ceicdata.com/en/india/foreign-exchange-rate-reserve-bank-of-india/foreign-exchange-rate-rbi-reference-rate-us-dollars>
13. CEIC. (2026b). India foreign exchange rate: RBI reference rate: pound sterling [FBIL reference rate, 17 September 2026]. CEIC Data. <https://www.ceicdata.com/en/india/foreign-exchange-rate-reserve-bank-of-india>
14. Christoffersen, P. F. (1998). Evaluating interval forecasts. *International Economic Review*, 39(4), 841–862.
15. Cowan, T., da Silva Rosa, R., & Walter, T. (n.d.). The demutualisation of the Australian Stock Exchange: Causes and consequences. University of Technology Sydney. <https://opus.lib.uts.edu.au/bitstream/10453/11608/1/2009002626.pdf>
16. CREDAI & ANAROCK. (2025, September 12). Indian REITs: A gateway to institutional real estate [Report presented at CREDAI Natcon]. CREDAI. https://credai.org/knowledge-center/credai-reports/reports/54_ANAROCK_Indian_REITs_CREDAI_Natcon_12_Sep_2025.pdf
17. Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007. <https://doi.org/10.2307/1912773>
18. Ferrell, A. (2007, July 12). A global view: Examining cross-border exchange mergers [Written testimony]. U.S. Senate Subcommittee on Securities, Insurance and Investment. <https://www.banking.senate.gov/imo/media/doc/ferrell.pdf>

19. Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2), 654–669.
<https://doi.org/10.1016/j.ejor.2017.11.054>
20. FSCS. (2019, April 1). FSCS announces higher protection limits of £85,000 for some products and services [Press release]. Financial Services Compensation Scheme.
<https://www.fscs.org.uk/media/press/2019/apr/higher-protection-limits/>
21. FSCS. (2026). What we cover. Financial Services Compensation Scheme.
<https://www.fscs.org.uk/what-we-cover>
22. Ghar.tv. (2025, November 25). India's REIT market records ₹2,331 crore distribution to over 3.3 lakh investors in Q2 FY26. Ghar.tv. https://www.ghar.tv/blog/blog_search.php?q=reit
23. Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5), 1779–1801.
24. GreenSutra. (2026, June 13). BRSR in 2026: All you need to know about SEBI's Business Responsibility and Sustainability Reporting. GreenSutra. <https://greensutra.in/news/brsr-all-you-need-to-know/>
25. Groww. (2026). NSE IPO: Check allotment status, listing date, listing price. Groww.
<https://groww.in/ipo/nse-ipo>
26. Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357–384.
27. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
28. Inclán, C., & Tiao, G. C. (1994). Use of cumulative sums of squares for retrospective detection of changes of variance. *Journal of the American Statistical Association*, 89(427), 913–923.
29. India TV News. (2026, September 24). NSE IPO listing highlights: Shares list at a premium of 0.84% over issue price, gain post-debut. India TV News. <https://www.indiatvnews.com/business/markets/nse-ipo-listing-price-today-on-bse-national-stock-exchange-nse-ipo-gmp-ahead-of-market-debut-nse-share-price-2026-09-24-1055115>
30. Indian Infrastructure. (2025, August 18). India's REIT market crosses Rs 1 trillion mark in market capitalisation. Indian Infrastructure. <https://indianinfrastructure.com/2025/08/18/indias-reit-market-crosses-rs-1-trillion-mark-in-market-capitalisation/>
31. kalilurrahman. (2026). NIFTY_50_STOCK_DATA [Data set]. GitHub.
https://github.com/kalilurrahman/NIFTY_50_STOCK_DATA
32. Keslio. (2026). Sustainability reporting requirements in India. Keslio.
<https://www.keslio.com/insights/sustainability-reporting-requirements-in-india>
33. KPMG. (2023, July). First Notes: SEBI framework on BRSR Core, value chain disclosures and assurance by listed entities. KPMG in India.
34. Kumar, D., & Maheswaran, S. (2012). Modelling asymmetry and persistence under the impact of sudden changes in the volatility of the Indian stock market. *IIMB Management Review*, 24(3).
<https://www.iimb.ac.in/imr/previous-issues/modelling-asymmetry-sep-2012.php>
35. Kupiec, P. H. (1995). Techniques for verifying the accuracy of risk measurement models. *The Journal of Derivatives*, 3(2), 73–84.
36. Lamoureux, C. G., & Lastrapes, W. D. (1990). Persistence in variance, structural change, and the GARCH model. *Journal of Business & Economic Statistics*, 8(2), 225–234.
37. Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1), 77–91.
<https://doi.org/10.1111/j.1540-6261.1952.tb01525.x>
38. Multibagg AI. (2026). NSE IPO co-location case: What changed now. Multibagg AI.
<https://www.multibagg.ai/market-pulse/articles/nse-ipo-co-location-case-update-cmu9ub3rc000c2zntfcyrgsmt>
39. National Stock Exchange of India. (2026a). Investor Protection Fund Trust: Compensation limit and corpus (corpus as on 31 May 2026, unaudited). NSE. <https://nseindia.com/complaints/investor-protection-fund-trust>
40. National Stock Exchange of India. (2026b, September 10). Red Herring Prospectus. NSE.
<https://www.nseindia.com/static/investor-relations/offer-documents>

41. Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59(2), 347–370.
42. News On AIR. (2024, August 14). Indian rupee closes at ₹83.94 against US dollar. News On AIR. <https://www.newsonair.gov.in/indian-rupee-closes-at-%e2%82%b983-94-against-us-dollar>
43. NYU Stern V-Lab. (2026). National Stock Exchange CNX Nifty index AGARCH volatility analysis (analysis last updated 20 July 2026). NYU Stern School of Business Volatility Laboratory. <https://vlab.stern.nyu.edu/volatility/VOL.NIFTY%3AIND-R.AGARCH>
44. OFX. (2024, August 14). GBP to INR exchange rates [Historical rate, 14 August 2024]. OFX. <https://ofx.com/en-ca/exchange-rates/gbp-to-inr>
45. Pandey, R., & Kumar, A. (2017). Modelling persistence in conditional volatility of asset returns. *Afro-Asian Journal of Finance and Accounting*, 7(1), 16–34.
46. Sansó, A., Aragón, V., & Carrion-i-Silvestre, J. L. (2004). Testing for changes in the unconditional variance of financial time series. *Revista de Economía Financiera*, 4, 32–53.
47. SEBI. (2021). Business responsibility and sustainability reporting by listed entities [Circular dated 10 May 2021; amendment to Regulation 34(2)(f) notified 5 May 2021, SEBI/LAD-NRO/GN/2021/22]. Securities and Exchange Board of India.
48. Share.market. (2026). How is the stock market regulated? Share. market. <https://www.share.market/buzz/learn/how-is-the-stock-market-regulated/>
49. Sinha, B. (2007). Modeling stock market volatility in emerging markets: Evidence from India. MPRA Paper No. 102455, Munich Personal RePEc Archive. <https://mpra.ub.uni-muenchen.de/102455/>
50. SIPC. (2026). What SIPC protects. Securities Investor Protection Corporation, United States. <https://www.sipc.org>
51. Steil, B. (2002). Changes in the ownership and governance of securities exchanges: Causes and consequences. *Brookings-Wharton Papers on Financial Services*, 2002(1), 61–91.
52. Suganya, D., Padmavathi, M., & Sarantinos, V. (2026). Volatility dynamics in Indian stock markets: Evidence from the post-2015 era. *Journal of Risk and Financial Management*, 19(7), 471. <https://doi.org/10.3390/jrfm19070471>
53. The Statesman. (2025). India's REIT market surpasses Hong Kong at Rs 2.3 lakh crore. The Statesman. <https://www.thestatesman.com/business/indias-reit-market-surpasses-hong-kong-at-rs-2-3-lakh-crore-1503529582.html>
54. U.S. Securities and Exchange Commission, Office of Investor Education and Advocacy. (n.d.). Investor bulletin: SIPC protection, part 1, SIPC basics. Investor.gov. <https://www.investor.gov/introduction-investing/general-resources/news-alerts/alerts-bulletins/investor-bulletins/investor-bulletin-sipc-protection-part-1-sipc-basics>
55. World Bank. (2025, December 17). Gross domestic product per capita, current US dollars: India, United Kingdom and United States (series PCAGDPINA646NWDB, PCAGDPGBA646NWDB and PCAGDPUSA646NWDB) [World Development Indicators, retrieved from FRED, Federal Reserve Bank of St. Louis]. <https://fred.stlouisfed.org/series/PCAGDPINA646NWDB>
56. Yahoo Finance. (n.d.). Knowledge Realty Trust (KRT.BO): Company profile. Yahoo Finance. <https://finance.yahoo.com/quote/KRT.BO/>